Infringement Claim Chart for US8265089B2 V. Cloudflare Inc ("Defendant")

Claims	Evidence
1. A computer communication network system comprising: a source computer, an MPDU aggregating module, a connection-based network, a gateway, a receiver-side connectionless network, and a receiver computer, wherein:	Defendant's company (Cloudflare Inc) provides the end-to-end internet whereby connectionless LAN autonomous systems connect to BGP WAN edge routers for the connection-based pathway and encapsulation framing across the WAN to mitigate DDoS. It distributes traffic (data unit) across the network and chooses a fast, efficient route to deliver the traffic. It offers a reverse proxy CDN service and data centers (network) which are connected to the rest of the world via multiple providers. Type text here 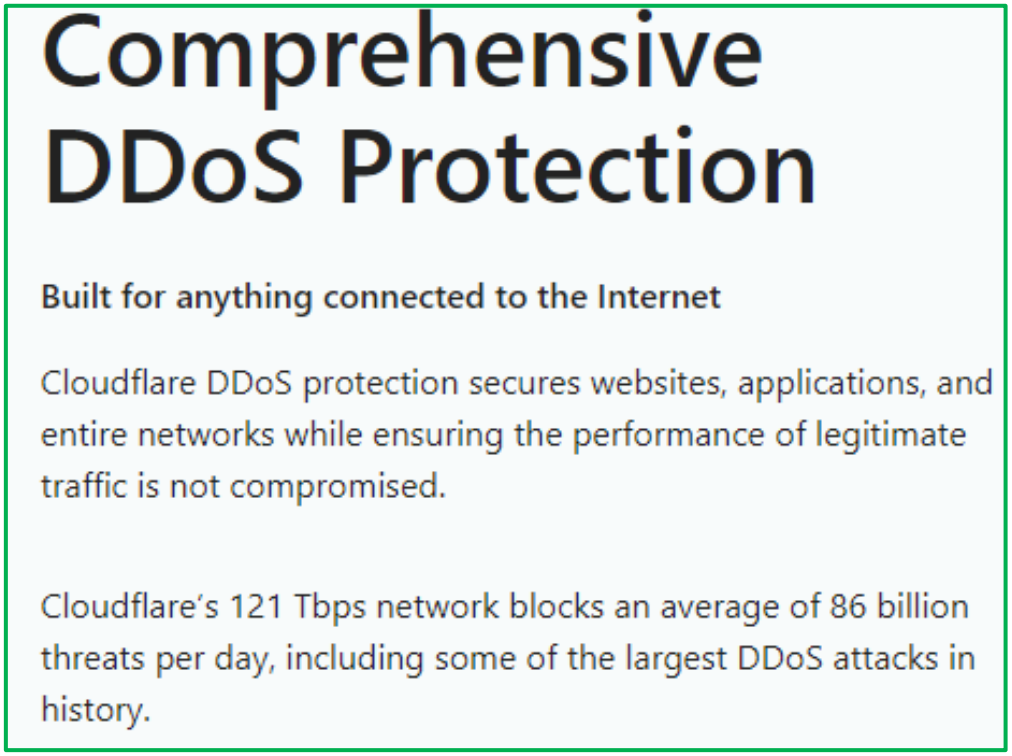 Source: https://www.cloudflare.com/ddos/

	Cloudflare, Inc. is an American web infrastructure and website security company that provides content delivery network and DDoS mitigation services.[2] Its services occur between a website's visitor and the Cloudflare customer's hosting provider, acting as a reverse proxy for websites.[3][4] Its headquarters are in San Francisco.[2] Source: https://en.wikipedia.org/wiki/Cloudflare#cite_note-CNBC-2 **Network:**CloudFlare's 23 data centers (internally we refer to them as PoPs) are connected to the rest of the world via multiple providers. These connections are both through transit (bandwidth) providers as well as other networks we directly peer with. Source: https://blog.cloudflare.com/cloudflares-architecture-eliminating-single-p/ 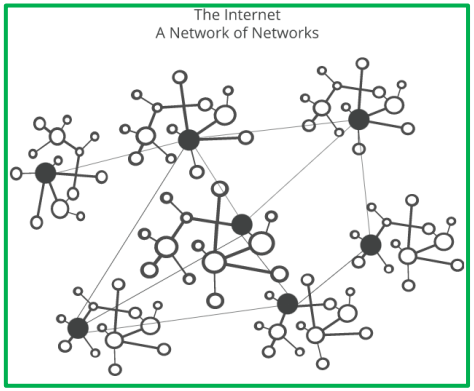 Source: https://www.cloudflare.com/en-in/learning/security/glossary/what-is-bgp/

Border Gateway Protocol (BGP) is the postal service of the Internet. When someone drops a letter into a mailbox, the Postal Service processes that piece of mail and chooses a fast, efficient route to deliver that letter to its recipient. Similarly, when someone submits data via the Internet, BGP is responsible for looking at all of the available paths that data could travel and picking the best route, which usually means hopping between autonomous systems.

BGP is the protocol that makes the Internet work by enabling data routing. When a user in Singapore loads a website with origin servers in Argentina, BGP is the protocol that enables that communication to happen quickly and efficiently.

Source: https://www.cloudflare.com/en-in/learning/security/glossary/what-is-bgp/

The Internet is a network of networks. It is broken up into hundreds of thousands of smaller networks known as autonomous systems (ASes). Each of these networks is essentially a large pool of routers run by a single organization.

Source: https://www.cloudflare.com/en-in/learning/security/glossary/what-is-bgp/

Source: https://www.cloudflare.com/network/

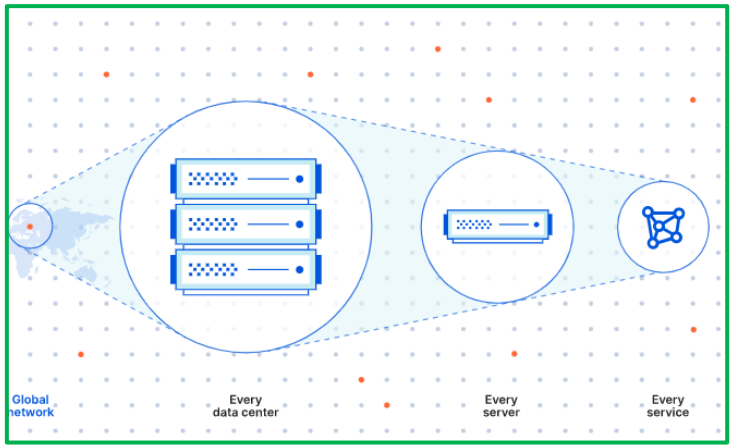

Source: https://www.cloudflare.com/network/

Anycast network

Minimize latency and increase resiliency with a global Anycast network that can effortlessly stop even the largest DDoS attacks.

Source: https://www.cloudflare.com/network/

Network administrators

Improve the performance, reliability, and security of your WAN by using the Cloudflare network.

Source: https://www.cloudflare.com/network/

	If your application requires the client IP and supports Proxy Protocol, enable **Proxy Protocols**. Proxy Protocol is a method for a proxy like Cloudflare to send the client IP to the origin application. Source: https://developers.cloudflare.com/spectrum/get-started/ The main goal of Gatebot was to automate as much of the mitigation workflow as possible. That means: to observe the network and note the anomalies, understand the targets of attacks and their metadata (such as the type of customer involved), and perform appropriate mitigation action. Source: https://blog.cloudflare.com/meet-gatebot-a-bot-that-allows-us-to-sleep/ **What is a reverse proxy? | Proxy servers explained** A reverse proxy protects web servers from attacks and can provide performance and reliability benefits. Learn more about forward and reverse proxies. Source: https://www.cloudflare.com/en-in/learning/cdn/glossary/reverse-proxy/

	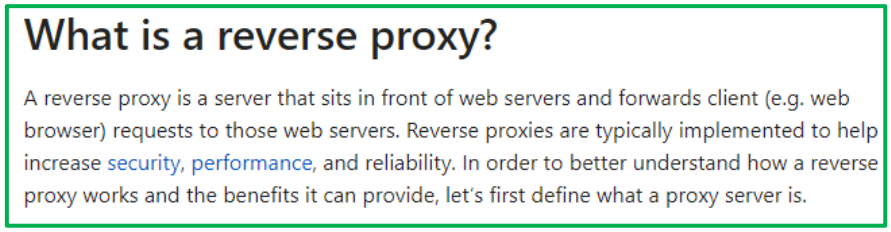 Source: https://www.cloudflare.com/en-in/learning/cdn/glossary/reverse-proxy/ 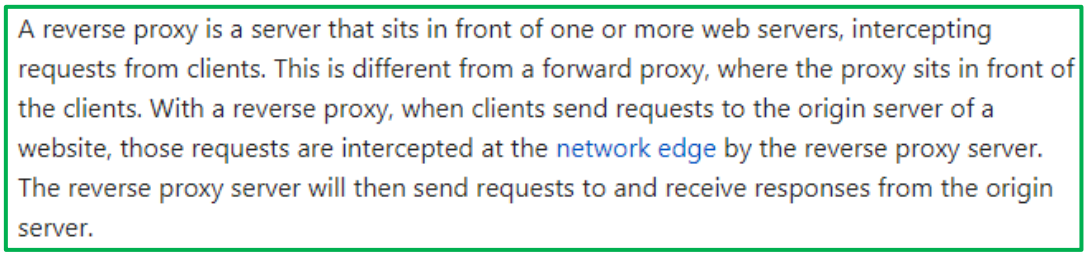 Source: https://www.cloudflare.com/en-in/learning/cdn/glossary/reverse-proxy/
the source computer is structured, and/or data-communication-connected to send a first packet, with the first packet including destination information indicating that it	The source computer is structured, and/or data-communication-connected to send the first packet, with the first packet including origin and destination information (IP addresses) indicating that it is intended to be sent to and received by the receiver computer. 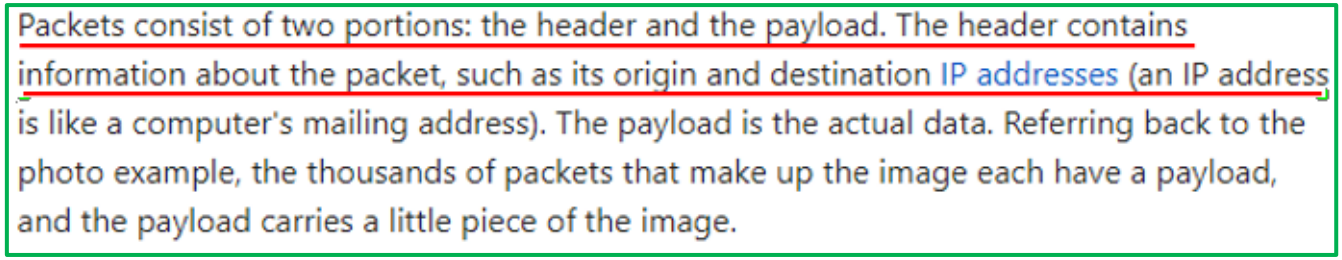 Source: https://www.cloudflare.com/en-in/learning/network-layer/what-is-a-packet/

is intended to be sent to and received by the receiver computer;	The Internet Protocol (IP) is a protocol, or set of rules, for routing and addressing packets of data so that they can travel across networks and arrive at the correct destination. Data traversing the Internet is divided into smaller pieces, called packets. IP information is attached to each packet, and this information helps routers to send packets to the right place. Every device or domain that connects to the Internet is assigned an IP address, and as packets are directed to the IP address attached to them, data arrives where it is needed. Source: https://www.cloudflare.com/learning/network-layer/internet-protocol/ 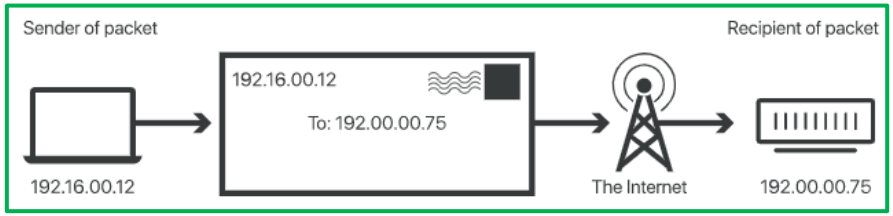 Source: https://www.cloudflare.com/learning/network-layer/internet-protocol/ Border Gateway Protocol (BGP) is the postal service of the Internet. When someone drops a letter into a mailbox, the Postal Service processes that piece of mail and chooses a fast, efficient route to deliver that letter to its recipient. Similarly, when someone submits data via the Internet, BGP is responsible for looking at all of the available paths that data could travel and picking the best route, which usually means hopping between autonomous systems. BGP is the protocol that makes the Internet work by enabling data routing. When a user in Singapore loads a website with origin servers in Argentina, BGP is the protocol that enables that communication to happen quickly and efficiently. Source: https://www.cloudflare.com/en-in/learning/security/glossary/what-is-bgp/

	The Internet is a network of networks. It is broken up into hundreds of thousands of smaller networks known as autonomous systems (ASes). Each of these networks is essentially a large pool of routers run by a single organization. Source: https://www.cloudflare.com/en-in/learning/security/glossary/what-is-bgp/ 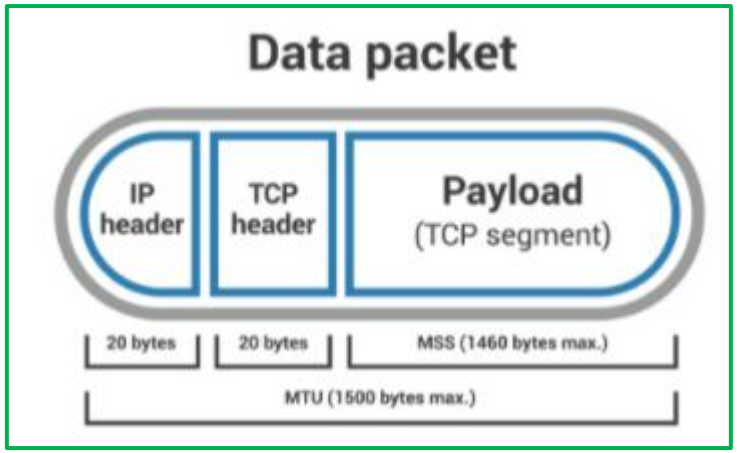 Source: https://www.cloudflare.com/en-in/learning/network-layer/what-is-mss/
the MPDU aggregating module is structured, programmed and/or data-communication-connected to receive the first	The defendant's product receives the first packet from the source computer and aggregates it into a first MPDU, where the first MPDU is in a form and format suitable to be communicated over the connection-based network. In networking, a packet is a small segment of a larger message. Data sent over computer networks*, such as the Internet, is divided into packets. These packets are then recombined by the computer or device that receives them. Source: https://www.cloudflare.com/en-in/learning/network-layer/what-is-a-packet/

packet from the source computer and to aggregate it into a first MPDU, where the first MPDU is in a form and format suitable to be communicated over the connection-based network;	MSS (maximum segment size) limits the size of packets, or small chunks of data, that travel across a network, such as the Internet. All data that travels over a network is broken up into packets. Packets have several headers attached to them that contain information about their contents and destination. MSS measures the non-header portion of a packet, which is called the payload. Source: https://www.cloudflare.com/en-in/learning/network-layer/what-is-mss/ 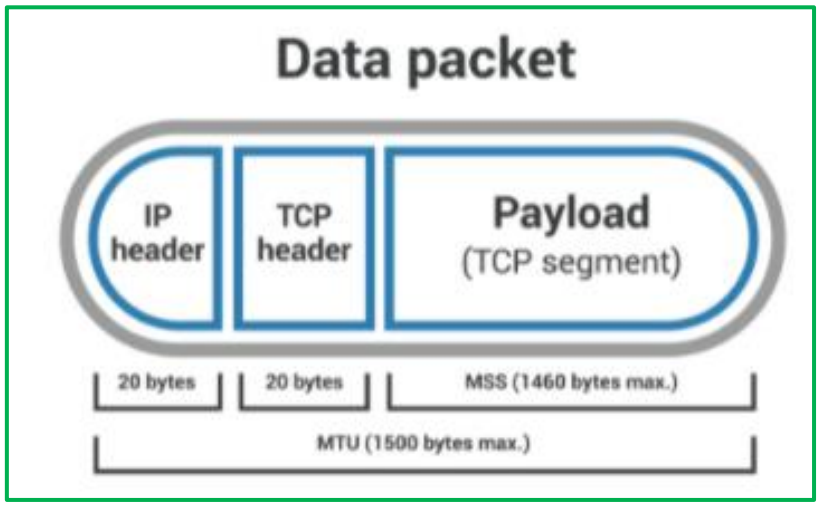 Source: https://www.cloudflare.com/en-in/learning/network-layer/what-is-mss/ The Internet is a network of networks. It is broken up into hundreds of thousands of smaller networks known as autonomous systems (ASes). Each of these networks is essentially a large pool of routers run by a single organization. Source: https://www.cloudflare.com/en-in/learning/security/glossary/what-is-bgp/

	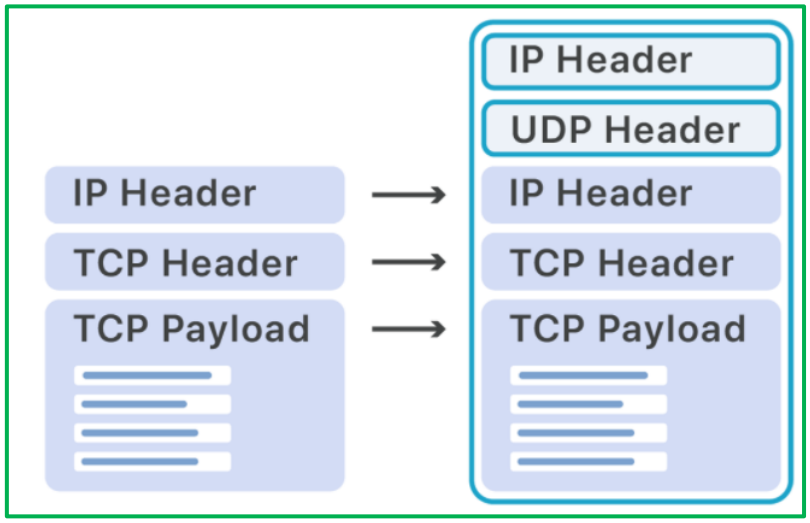 Source: https://blog.cloudflare.com/high-availability-load-balancers-with-maglev/ It's possible for routers to terminate BGP sessions after a much shorter delay using the Bidirectional Forwarding Detection (BFD) protocol between the router and load balancers. Different routers have different limitations and restrictions on BFD that makes it difficult to use in an environment heavily using L2 link aggregation and VXLANs. Source: https://blog.cloudflare.com/high-availability-load-balancers-with-maglev/
the connection-based network is structured, programmed and/or data-	CloudFlare uses its own DNS to point the BGP WAN paths to its Gateways via WAN ANYCAST within the WAN Gateways (over 150 distributed around the world at this point). By announcing these customer IP addresses it now sets up for a remote gateway intercept as the normal connection-based path to the connectionless network (AS) of the intended receiving computer is “hijacked” by CloudFlare.

communication-connected to receive the first MPDU from the MPDU aggregating module and to communicate it to the gateway in a connection-based manner;	The cause of the outage was a system-wide failure of our edge routers. CloudFlare currently runs 23 data centers worldwide. These data centers are connected to the rest of the Internet using routers. These routers announce the path that, from any point on the Internet, packets should use to reach our network. When a router goes down, the routes to the network that sits behind the router are withdrawn from the rest of the Internet. Source: https://blog.cloudflare.com/todays-outage-post-mortem-82515/ At CloudFlare, we use Anycast at two levels: the WAN and the LAN. At the WAN level, every router in all of CloudFlare's 23 data centers announces all of our external-facing IP addresses. For example, one of the IPs that CloudFlare announces for DNS services is 173.245.58.205. A route to that IP address is announced from all 23 CloudFlare data centers. When you send a packet to that IP address, it passes through a series of routers. Those routers look at the available paths to CloudFlare's end points and send the packet down the one with the fewest stops along the way (i.e., "hops"). You can run a traceroute to see each of these steps. Source: https://blog.cloudflare.com/cloudflares-architecture-eliminating-single-p/ 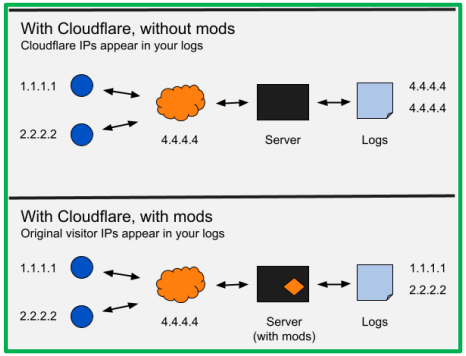 Source: https://support.cloudflare.com/hc/en-us/articles/200170786-How-do-I-restore-original-visitor-IP-with-Nginx-

the gateway is structured, programmed and/or data-communication-connected to receive the first MPDU from the connection-based network, to disaggregate the first MPDU into a plurality of smaller data units (DUs) including a first DU at least partially constituted by the first packet, and to selectively communicate the first DU to the receiver-side connectionless network;	The components of CloudFlare's edge systems are deployed at the edge of the network. The gateway receives the first MPDU and disaggregates the first MPDU into a plurality of smaller data units (DUs) including a first DU at least partially constituted by the first packet and selectively communicates the first DU to the receiver-side connectionless network. The Cloudflare Network is acting as the Gateway, and it is in data communication to the receiver computer. **Flowspec** We are largely a Juniper shop at CloudFlare and all the edge routers that were affected were from Juniper. One of the reasons we like Juniper is their support of a protocol called Flowspec. Flowspec allows you to propagate router rules to a large number of routers efficiently. At CloudFlare, we constantly make updates to the rules on our routers. We do this to fight attacks as well as to shift traffic so it can be served as fast as possible. Source: https://blog.cloudflare.com/todays-outage-post-mortem-82515/ 1. **Network:**CloudFlare's 23 data centers (internally we refer to them as PoPs) are connected to the rest of the world via multiple providers. These connections are both through transit (bandwidth) providers as well as other networks we directly peer with. 2. **Router:**at the edge of each of our PoPs is a router. This router announces the paths packets take to CloudFlare's network from the rest of the Internet. 3. **Switch:** within each PoP there will be one or more switches that aggregate traffic within the PoP's local area network (LAN). 4. **Server:** behind each switch there are a collection of servers. These servers perform some of the key tasks to power CloudFlare's service including DNS resolution, proxying, caching, and logging. Source: https://blog.cloudflare.com/cloudflares-architecture-eliminating-single-p/

	Those are the four components you'll find in the racks that we run in locations around the world. You'll notice some things from a typical hardware stack seem to be missing. For example, there's no hardware load balancer. The problem with hardware load balancers (and hardware firewalls, for that matter) is that they often become the bottleneck and create a single point of failure themselves. Instead of relying on a piece of hardware to load balance across our network, we use routing protocols to spread traffic and handle failure. Source: https://blog.cloudflare.com/cloudflares-architecture-eliminating-single-p/
the receiver-side connectionless network is structured, programmed and/or data-communication-connected to receive the first DU from the gateway on condition that it was selectively communicated by the gateway, and to	All traffic goes through the CloudFlare network and then is proxied back to the origin servers. The CloudFlare Network is acting as the gateway and it is in data communication to the receiver computer. What is a reverse proxy? | Proxy servers explained A reverse proxy protects web servers from attacks and can provide performance and reliability benefits. Learn more about forward and reverse proxies. Source: https://www.cloudflare.com/en-in/learning/cdn/glossary/reverse-proxy/

communicate at least the first data packet portion of the first DU to the receiver computer in a connectionless manner;	**What is a reverse proxy?** A reverse proxy is a server that sits in front of web servers and forwards client (e.g. web browser) requests to those web servers. Reverse proxies are typically implemented to help increase security, performance, and reliability. In order to better understand how a reverse proxy works and the benefits it can provide, let's first define what a proxy server is. Source: https://www.cloudflare.com/en-in/learning/cdn/glossary/reverse-proxy/ A reverse proxy is a server that sits in front of one or more web servers, intercepting requests from clients. This is different from a forward proxy, where the proxy sits in front of the clients. With a reverse proxy, when clients send requests to the origin server of a website, those requests are intercepted at the network edge by the reverse proxy server. The reverse proxy server will then send requests to and receive responses from the origin server. Source: https://www.cloudflare.com/en-in/learning/cdn/glossary/reverse-proxy/
the gateway is structured, programmed and/or data-communication-connected to collect selected	CloudFlare's PoP (Gateway) peers to the connection-based WAN with BGP announcements that establish paths. To determine how it is "connected" the MPDU of the connecting BGP peer may include in its protocol the autonomous system number, the types of flow (NetFlow or Sflow as an example), carrier ethernet type, MPLS, etc. WAN is a mix of transport technologies. The internal tool that profiles attacks and outputs signatures that automated systems, as well as the ops team, can use to stop attacks. Often, these signatures create router rules to either rate limit or drop known-bad requests.

network protocol data from the first MPDU, with the selected network protocol data including at least some network protocol data included in the first MPDU and not included in any of the plurality of DUs, and with the selected network protocol data relating exclusively to network protocols and including no data from any data payload(s) which may be present in the MPDU;	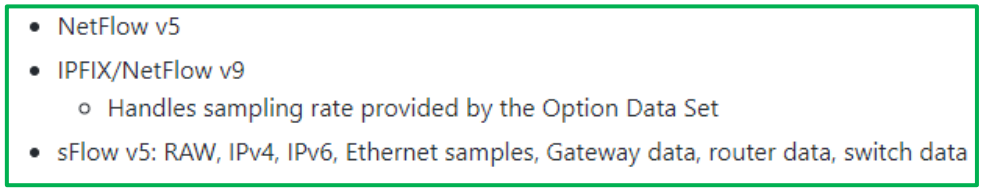 Source: https://github.com/cloudflare/goflow  Source: https://github.com/cloudflare/goflow 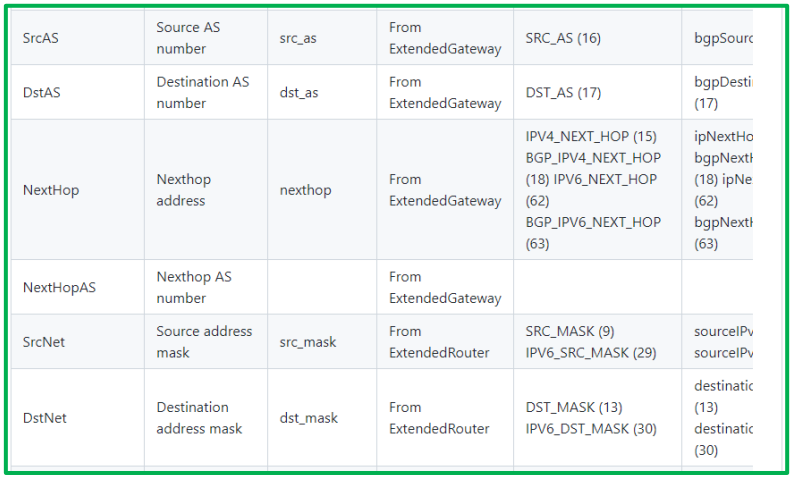 Source: https://github.com/cloudflare/goflow

SrcPort	Source port (when UDP/TCP/SCTP)	srcport	Included	L4_SRC_PORT (7)	sourceTra
DstPort	Destination port (when UDP/TCP/SCTP)	dstport	Included	L4_DST_PORT (11)	destinatic (11)

SrcAS	Source AS number	src_as	From ExtendedGateway	SRC_AS (16)	bgpSourc
DstAS	Destination AS number	dst_as	From ExtendedGateway	DST_AS (17)	bgpDesti (17)
NextHop	Nexthop address	nexthop	From ExtendedGateway	IPV4_NEXT_HOP (15) BGP_IPV4_NEXT_HOP (18) IPV6_NEXT_HOP (62) BGP_IPV6_NEXT_HOP (63)	ipNextHo bgpNextH (18) ipNe (62) bgpNextH (63)
NextHopAS	Nexthop AS number		From ExtendedGateway		
SrcNet	Source address mask	src_mask	From ExtendedRouter	SRC_MASK (9) IPV6_SRC_MASK (29)	sourceIPv sourceIPv
DstNet	Destination address mask	dst_mask	From ExtendedRouter	DST_MASK (13) IPV6_DST_MASK (30)	destinatic (13) destinatic (30)

SrcMac	Source mac address		Included	IN_SRC_MAC (56)	sourceMa
DstMac	Destination mac address		Included	OUT_DST_MAC (57)	postDesti (57)

Source: https://github.com/cloudflare/goflow

Through this website IP to ASN lookup,I can get a series of results about the ASN,but I can't find the meaning of BGP prefix.Maybe I know what it is,however,how it works?

Source: https://github.com/cloudflare/goflow

A prefix announced in BGP consists of the IPv4 or IPv6 address block being announced and also a path of AS numbers, indicating which ASNs the traffic must pass through to reach the announced address block.

Source: https://github.com/cloudflare/goflow

GoFlow

This application is a NetFlow/IPFIX/sFlow collector in Go.

It gathers network information (IP, interfaces, routers) from different flow protocols, serializes it in a protobuf format and sends the messages to Kafka using Sarama's library.

Source: https://github.com/cloudflare/goflow

How is it used at Cloudflare

The samples flowing into Kafka are **processed** and special fields are inserted using other databases:

- User plan
- Country
- ASN and BGP information

The extended protobuf has the same base of the one in this repo. The **compatibility** with other software is preserved when adding new fields (thus the fields will be lost if re-serialized).

Source: https://github.com/cloudflare/goflow

The pipeline at Cloudflare is connecting collectors with flow processors that will add more information: with IP address, add country, ASN, etc.

Source: https://github.com/cloudflare/goflow

The BGP information provided by routers can be unreliable (if the router does not have a BGP full-table or it is a static route). You can use Maxmind prefix to ASN in order to solve this issue.

Source: https://github.com/cloudflare/goflow

Because IBGP supports multihop connections, IBGP neighbors can be located anywhere within the autonomous system (AS) and often do not share a link. A recursive route lookup resolves the loopback peer address to an IP forwarding next hop. In this example, this service is provided by OSPF. Although interior gateway protocol (IGP) neighbors do not need to be directly connected, they do need to be fully meshed. In this case, fully meshed means that each device is logically connected to every other device through neighbor peer relationships. The `neighbor` statement creates the mesh.

Source:https://www.juniper.net/documentation/us/en/software/junos/bgp/topics/topic-map/bgp-peering-sessions.html#id-understanding-external-bgp-peering-sessions

	To establish point-to-point connections between peer autonomous systems (ASs), you configure a BGP session on each interface of a point-to-point link. Generally, such sessions are made at network exit points with neighboring hosts outside the AS. Figure 1 shows an example of a BGP peering session. Source:https://www.juniper.net/documentation/us/en/software/junos/bgp/topics/topic-map/bgp-peering-sessions.html#id-understanding-external-bgp-peering-sessions In Figure 1, Router A is a gateway router for AS 3, and Router B is a gateway router for AS 10. For traffic internal to either AS, an interior gateway protocol (IGP) is used (OSPF, for instance). To route traffic between peer ASs, a BGP session is used. You arrange BGP routing devices into groups of peers. Different peer groups can have different group types, AS numbers, and route reflector cluster identifiers. To define a BGP group that recognizes only the specified BGP systems as peers, statically configure all the system's peers by including one or more `neighbor` statements. The peer neighbor's address can be either an IPv6 or IPv4 address. Source:https://www.juniper.net/documentation/us/en/software/junos/bgp/topics/topic-map/bgp-peering-sessions.html#id-understanding-external-bgp-peering-sessions **What is a reverse proxy?** A reverse proxy is a server that sits in front of web servers and forwards client (e.g. web browser) requests to those web servers. Reverse proxies are typically implemented to help increase security, performance, and reliability. In order to better understand how a reverse proxy works and the benefits it can provide, let's first define what a proxy server is. Source: https://www.cloudflare.com/en-in/learning/cdn/glossary/reverse-proxy/

	A reverse proxy is a server that sits in front of one or more web servers, intercepting requests from clients. This is different from a forward proxy, where the proxy sits in front of the clients. With a reverse proxy, when clients send requests to the origin server of a website, those requests are intercepted at the network edge by the reverse proxy server. The reverse proxy server will then send requests to and receive responses from the origin server. Source: https://www.cloudflare.com/en-in/learning/cdn/glossary/reverse-proxy/
the gateway is further structured, programmed and/or data-communication-connected to apply a first rule to the selected network protocol data that has been collected by the gateway; and	Cloudflare has an internal tool that profiles attacks and outputs signatures that our automated systems as well as our ops team can use to stop attacks. Often, the signatures are used in order to create router rules to either rate limit or drop known-bad requests. Cloudflare's DDoS mitigation team has developed a solution based on kernel bypass and classic BPF. This allows filtering network packets in userspace, skipping the usual packet. Sometimes though, even this is not sufficient. Iptables is fast in the general case, but has its limits. During very large attacks, exceeding 1M packets per second per server, we shift the attack traffic from kernel iptables to a kernel bypass user space program (which we call floodgate). We use a partial kernel bypass solution using Solarflare EF_VI interface. With this on each server we can process more than 5M attack packets per second while consuming only a single CPU core. With floodgate we have comfortable amount of CPU left for our applications, even during the largest network events. Source: https://blog.cloudflare.com/meet-gatebot-a-bot-that-allows-us-to-sleep/

	Cloudflare DDoS protection secures websites, applications, and entire networks while ensuring the performance of legitimate traffic is not compromised. Cloudflare's 121 Tbps network blocks an average of 86 billion threats per day, including some of the largest DDoS attacks in history. Source: https://www.cloudflare.com/ddos/ Cloudflare Spectrum is a reverse proxy service that provides DDoS protection for any application (not just the web), such as FTP, SSH, VoIP, gaming, or any application running over a TCP/UDP protocol. Spectrum comes with built-in load balancing and traffic acceleration for L4 traffic. Source: https://www.cloudflare.com/ddos/

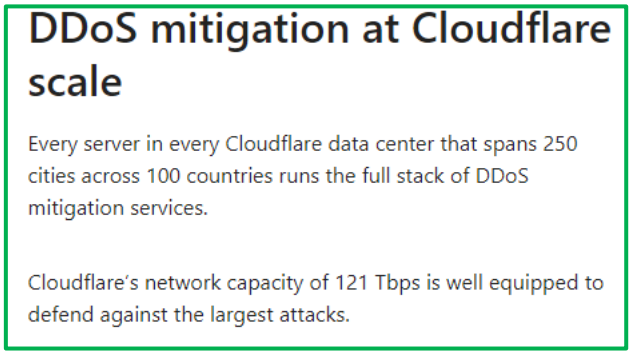

Source: https://www.cloudflare.com/ddos/

> Finally, there are a number of tweaks we can make on at the HTTP layer. For specific attacks we disable HTTP Keep-Alives forcing attackers to re-establish TCP sessions for each request. This sacrifices a bit of performance for valid traffic as well, but is a surprisingly powerful tool throttling many attacks. For other attack patterns we turn the "I'm under attack" mode on, forcing the attack to hit our JavaScript challenge page.

Source: https://blog.cloudflare.com/meet-gatebot-a-bot-that-allows-us-to-sleep/

> During normal operations our attitude to attacks is rather pragmatic. Since the inbound traffic is distributed across hundreds of servers we can survive periodic spikes and small attacks without doing anything. Vanilla Linux is remarkably resilient against unexpected network events. This is especially true since kernel 4.4 when the performance of SYN cookies was greatly improved.

Source: https://blog.cloudflare.com/meet-gatebot-a-bot-that-allows-us-to-sleep/

	Gatebot is much faster and much more precise than even our most experienced SREs. Without Gatebot we wouldn't be able to operate our service with the appropriate level of confidence. Furthermore, Gatebot has proved to be remarkably adaptable - we started by automating handling of Layer 3 attacks, but soon we proved that the general model works well for automating other things. Today we have more than 10 separate Gatebot instances doing everything from mitigating Layer 7 attacks to informing our Customer Support team of misbehaving customer origin servers. Source: https://blog.cloudflare.com/meet-gatebot-a-bot-that-allows-us-to-sleep/ When building a DDoS mitigation service it's incredibly tempting to think that the solution is scrubbing centers or scrubbing servers. I, too, thought that was a good idea in the beginning, but experience has shown that there are serious pitfalls to this approach.... Source: https://blog.cloudflare.com/meet-gatebot-a-bot-that-allows-us-to-sleep/
the gateway is further structured, programmed and/or data-	Cloudflare has an internal tool that profiles attacks and outputs signatures that our automated systems as well as our ops team can use to stop attacks. Often, the signatures are used in order to create router rules to either rate limit or drop known-bad requests. Cloudflare's DDoS mitigation team has developed a solution based on kernel bypass and classic BPF. This allows filtering network packets in userspace, skipping the usual packet.

communication-connected to selectively make a responsive reaction based, at least in part, upon the application of the first rule applied by the gateway to the selected network protocol data.	Gatebot is much faster and much more precise than even our most experienced SREs. Without Gatebot we wouldn't be able to operate our service with the appropriate level of confidence. Furthermore, Gatebot has proved to be remarkably adaptable - we started by automating handling of Layer 3 attacks, but soon we proved that the general model works well for automating other things. Today we have more than 10 separate Gatebot instances doing everything from mitigating Layer 7 attacks to informing our Customer Support team of misbehaving customer origin servers. Source: https://blog.cloudflare.com/meet-gatebot-a-bot-that-allows-us-to-sleep/ During normal operations our attitude to attacks is rather pragmatic. Since the inbound traffic is distributed across hundreds of servers we can survive periodic spikes and small attacks without doing anything. Vanilla Linux is remarkably resilient against unexpected network events. This is especially true since kernel 4.4 when the performance of SYN cookies was greatly improved. Source: https://blog.cloudflare.com/meet-gatebot-a-bot-that-allows-us-to-sleep/ Gilberto Bertin discusses the architecture of Cloudflare's automatic DDoS mitigation pipeline, the initial packet filtering solution based on Iptables, and why Cloudflare had to introduce userspace offload. Bertin also describes how they switched from a proprietary offload technology to XDP for network stack bypass and how they are using XDP to load balance traffic. Source: https://www.infoq.com/presentations/xdp-ddos-cloudflare/

	 Source: https://www.cloudflare.com/ddos/